P⊛RTAL
USPTO

Search: ⦿ The ACM Digital Library   ◯ The Guide

|+database +checkpoint +switch | **SEARCH**

L... C . . BRARY

● Feedback Report a problem Satisfaction survey

Terms used **database checkpoint switch**   Found **349** of **169,166**

Sort results by | relevance | ● Save results to a Binder   Try an Advanced Search
Display results | expanded form | ? Search Tips   Try this search in The ACM Guide
                                           ☐ Open results in a new window

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10   next
Best 200 shown   Relevance scale ☐▭▬■■

1 Performance measurement methodologies for database systems   ■
Steven A. Demurjian, David K. Hsiao, Douglas S. Kerr, Robert C. Tekampe, Robert J. Watson
October 1985 **Proceedings of the 1985 ACM annual conference on The range of computing : mid-80's perspective: mid-80's perspective**
Publisher: ACM Press
Full text available: 📄 pdf(1.34 MB)   Additional Information: full citation, references, index terms

2 An implementation for small databases with high availability   ■
Kwok-yan Lam
October 1991 **ACM SIGOPS Operating Systems Review,** Volume 25 Issue 4
Publisher: ACM Press
Full text available: 📄 pdf(610.33 KB)   Additional Information: full citation, abstract, index terms

A technique for implementing the kind of small databases that frequently occur in the design of operating systems and distributed systems is presented. This technique is based on the checkpointing mechanism for maintaining data persistence. It can be regarded as a hybrid of the stable storage approach and the master-slave approach to fault tolerance, and is especially suitable for systems that support the processor pool paradigm of computation. This scheme has been used to implement a monitoring ...

3 Database systems management and Oracle8   ■
C. Gregory Doherty
June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98,** Volume 27 Issue 2
Publisher: ACM Press
Full text available: 📄 pdf(290.86 KB)   Additional Information: full citation, abstract, references, index terms

Oracle's corporate mission is to enable the Information Age through network computing, a vision of broader access to information for all and the empowerment and increased productivity that can result. The technology implications of the network computing vision are ubiquitous access via low-cost appliances to smaller numbers of larger databases, accessed via professionally managed networks compliant with open internetworking protocols. The latest release of the Oracle data s ...

4 SafetyNet: improving the availability of shared memory multiprocessors with global checkpoint/recovery   ■

Daniel J. Sorin, Milo M. K. Martin, Mark D. Hill, David A. Wood
May 2002 **ACM SIGARCH Computer Architecture News , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02 , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02**, Volume 30 Issue 2
Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(1.28 MB) Publisher Site   Additional Information: full citation, abstract, references, citings, index terms

We develop an availability solution, called *SafetyNet,* that uses a unified, lightweight checkpoint/recovery mechanism to support multiple long-latency fault detection schemes. At an abstract level, *SafetyNet* logically maintains multiple, globally consistent checkpoints of the state of a shared memory multiprocessor (i.e., processors, memory, and coherence permissions), and it recovers to a pre-fault checkpoint of the system and re-executes if a fault is detected. *SafetyNet* e ...

Keywords: availability, shared memory, multiprocessor

5 Increasing cache port efficiency for dynamic superscalar microprocessors

Kenneth M. Wilson, Kunle Olukotun, Mendel Rosenblum
May 1996 **ACM SIGARCH Computer Architecture News , Proceedings of the 23rd annual international symposium on Computer architecture ISCA '96**, Volume 24 Issue 2
Publisher: ACM Press

Full text available: pdf(1.09 MB)   Additional Information: full citation, abstract, references, citings, index terms

The memory bandwidth demands of modern microprocessors require the use of a multi-ported cache to achieve peak performance. However, multi-ported caches are costly to implement. In this paper we propose techniques for improving the bandwidth of a single cache port by using additional buffering in the processor, and by taking maximum advantage of a wider cache port. We evaluate these techniques using realistic applications that include the operating system. Our techniques using a single-ported ca ...

6 Miscellaneous: Wanted: an application aware checkpointing service

Colin Allison
September 1994 **Proceedings of the 6th workshop on ACM SIGOPS European workshop: Matching operating systems to application needs**
Publisher: ACM Press

Full text available: pdf(506.82 KB)   Additional Information: full citation, abstract, references

Checkpointing is an essential component of many application frameworks but OS kernels rarely offer assistance with this task although they are in a very good position to do so. OS process management could and should be extended to accommodate this common need and provide an application aware system service.

7 Recovery Techniques for Database Systems

Joost S. M. Verhofstad
June 1978 **ACM Computing Surveys (CSUR)**, Volume 10 Issue 2
Publisher: ACM Press

Full text available: pdf(2.32 MB)   Additional Information: full citation, references, citings, index terms

8 MMDB reload algorithms

Le Gruenwald, Margaret H. Eich

April 1991 **ACM SIGMOD Record , Proceedings of the 1991 ACM SIGMOD international conference on Management of data SIGMOD '91**, Volume 20 Issue 2
Publisher: ACM Press
Full text available: pdf(984.01 KB)   Additional Information: <u>full citation</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>

---

9 <u>On-line reorganization of sparsely-populated B+-trees</u>
Chendong Zou, Betty Salzberg
June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international conference on Management of data SIGMOD '96**, Volume 25 Issue 2
Publisher: ACM Press
Full text available: pdf(1.17 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>

In this paper, we present an efficient method to do online reorganization of sparsely-populated B+-trees. It reorganizes the leaves first, compacting in short operations groups of leaves with the same parent. After compacting, optionally, the new leaves may swap locations or be moved into empty pages so that they are in key order on the disk. After the leaves are reorganized, the method shrinks the tree by making a copy of the upper part of the tree while leaving the leaves in place. ...

---

10 <u>Highly available systems for database applications</u>
Won Kim
March 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 1
Publisher: ACM Press
Full text available: pdf(2.43 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>, <u>review</u>

As users entrust more and more of their applications to computer systems, the need for systems that are continuously operational (24 hours per day) has become even greater. This paper presents a survey and analysis of representative architectures and techniques that have been developed for constructing highly available systems for database applications. It then proposes a design of a distributed software subsystem that can serve as a unified framework for constructing database applica ...

---

11 <u>50,000 users on an Oracle8 universal server database</u>
Tirthankar Lahiri, Ashok Joshi, Amit Jasuja, Sumanta Chatterjee
June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98**, Volume 27 Issue 2
Publisher: ACM Press
Full text available: pdf(460.50 KB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>

In this paper, we describe the Oracle Large User Population Demonstration and highlight the scalability mechanisms in the Oracle8 Universal Data Server which make it possible to support as many as 50,000 concurrent users on a single Oracle8 database without any middle-tier TP-monitor software. Supporting such large user populations requires many mechanisms for high concurrency and throughput. Algorithms in all areas of the server ranging from process and buffer management to SQL compilation ...

---

12 <u>A simple and efficient implementation of a small database</u>
A. Birrell, M. Jones, E. Wobber
November 1987 **ACM SIGOPS Operating Systems Review , Proceedings of the eleventh ACM Symposium on Operating systems principles SOSP '87**, Volume 21 Issue 5
Publisher: ACM Press

Full text available: pdf(689.45 KB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>

This paper describes a technique for implementing the sort of small databases that frequently occur in the design of operating systems and distributed systems. We take advantage of the existence of very large virtual memories, and quite large real memories, to make the technique feasible. We maintain the database as a strongly typed data structure in virtual memory, record updates incrementally on disk in a log and occasionally make a checkpoint of the entire database. We recover fr ...

13 Multiprocessor main memory transaction processing

K. Li, J. F. Naughton
January 2000 **Proceedings of the first international symposium on Databases in parallel and distributed systems**
Publisher: IEEE Computer Society Press

Full text available: pdf(1.16 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>

In this paper we describe an experiment designed to evaluate the potential transaction processing system performance achievable through the combination of multiple processors and massive memories. The experiment consisted of the design and implementation of a transaction processing kernel on stock multiprocessors. We found that with sufficient memory, multiple processors can greatly improve performance. A prototype implementation of the kernel on a pair of Firefly multiprocessors (each with ...

14 Physical integrity in a large segmented database

Raymond A. Lorie
March 1977 **ACM Transactions on Database Systems (TODS)**, Volume 2 Issue 1
Publisher: ACM Press

Full text available: pdf(1.12 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>

A database system can generally be divided into three major components. One component supports the logical database as seen by the user. Another component maps the information into physical records. The third component, called the storage component, is responsible for mapping these records onto auxiliary storage (generally disks) and controlling their transfer to and from main storage. This paper is primarily concerned with the implementation of a storage component. It considers ...

Keywords: checkpoint-restart, database, recovery, storage management

15 A reusable lightweight executive for command and control systems

Nathan Fleener, Laura Moody, Mary Stewart
November 1998 **ACM SIGAda Ada Letters , Proceedings of the 1998 annual ACM SIGAda international conference on Ada SIGAda '98**, Volume XVIII Issue 6
Publisher: ACM Press
Full text available: pdf(678.14 KB)   Additional Information: <u>full citation</u>, <u>references</u>, <u>index terms</u>

Keywords: Ada, executive, portability

16 A database cache for high performance and fast restart in database systems

Klaus Elhardt, Rudolf Bayer
December 1984 **ACM Transactions on Database Systems (TODS)**, Volume 9 Issue 4

Publisher: ACM Press

Full text available: pdf(1.72 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>, <u>review</u>

Performance in database systems is strongly influenced by buffer management and transaction recovery methods. This paper presents the principles of the database cache, which replaces the traditional buffer. In comparison to buffer management, cache management is more carefully coordinated with transaction management, and integrates transaction recovery. High throughput of small- and medium-sized transactions is achieved by fast commit processing and low database traffic. Very fas ...

17 <u>Data replicas in distributed information services</u>

H. M. Gladney

March 1989 **ACM Transactions on Database Systems (TODS)**, Volume 14 Issue 1

Publisher: ACM Press

Full text available: pdf(1.94 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>index terms</u>, <u>review</u>

In an information distribution network in which records are repeatedly read, it is cost-effective to keep read-only copies in work locations. This paper presents a method of updating replicas that need not be immediately synchronized with the source data or with each other. The method allows an arbitrary mapping from source records to replica records. It is fail-safe, maximizes workstation autonomy, and is well suited to a network with slow, unreliable, and/or expensive communications links ...

18 <u>Integrating association rule mining with relational database systems: alternatives and</u> <u>implications</u>

Sunita Sarawagi, Shiby Thomas, Rakesh Agrawal

June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98**, Volume 27 Issue 2

Publisher: ACM Press

Full text available: pdf(2.03 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>

Data mining on large data warehouses is becoming increasingly important. In support of this trend, we consider a spectrum of architectural alternatives for coupling mining with database systems. These alternatives include: loose-coupling through a SQL cursor interface; encapsulation of a mining algorithm in a stored procedure; caching the data to a file system on-the-fly and mining; tight-coupling using primarily user-defined functions; and SQL implementations for processing in the DBMS. We ...

19 <u>Novel ideas: Using variable-MHz microprocessors to efficiently handle uncertainty in</u> <u>real-time systems</u>

Eric Rotenberg

December 2001 **Proceedings of the 34th annual ACM/IEEE international symposium on Microarchitecture**

Publisher: IEEE Computer Society

Full text available: pdf(1.06 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>
<u>Publisher Site</u>

Guaranteed performance is critical in real-time systems because correct operation requires tasks complete on time. Meanwhile, as software complexity increases and deadlines tighten, embedded processors inherit high-performance techniques such as pipelining, caches, and branch prediction. Guaranteeing the performance of complex pipelines is difficult and worst-case analysis often under-estimates the microarchitecture for correctness. Ultimately, the designer must turn to clock frequency as a reli ...

20 Design and implementation of a portable database system for small computers

G. Weck, B. Wiesner

October 1981 **Proceedings of the 1981 ACM SIGSMALL symposium on Small systems and SIGMOD workshop on Small database systems**

Publisher: ACM Press

Full text available: pdf(732.79 KB)   Additional Information: full citation, abstract, references, index terms

This paper describes the design and implementation of PEDMS. Part B describes design objectives, resulting system architecture and interfaces to user and database administrator. Part C concentrates on the implementation of PEDMS, discussing the difficulties imposed by the environment, and the ways of circumventing these difficulties.

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10   next

Useful downloads: Adobe Acrobat   QuickTime   Windows Media Player   Real Player